AI Governance Business Case Template - Tractatus Framework

Tractatus AI Safety Framework

2025-10-25

Executive Summary

Status: [DRAFT - REQUIRES COMPLETION WITH ORGANIZATIONAL DATA]

Current AI Governance Posture

- **Current AI systems deployed:** [NUMBER] systems across [NUMBER] departments
- **Regulatory exposure:** [List applicable regulations: EU AI Act, sector-specific, etc.]
- **Known governance gaps:** [List identified gaps from current state assessment]
- **Risk appetite:** [Conservative / Moderate / Aggressive]

Proposed Approach: Tractatus Framework

The Tractatus Framework is a **research/development framework** for AI governance that uses architectural controls to manage AI decision boundaries. It is designed to help organizations:

- Define which decisions require human approval
- Maintain instruction persistence across AI sessions
- Monitor AI system behavior under operational pressure
- Create audit trails for compliance purposes

Framework Status: Early-stage research implementation. Organizations should evaluate readiness for adapting research frameworks vs. waiting for mature commercial solutions.

Decision Required

- **Investment:** [ESTIMATED COST - requires vendor engagement]
- **Timeline:** [PROJECTED TIMELINE - depends on organizational complexity]
- **Alternatives considered:** [List other approaches evaluated]
- **Recommendation:** [PENDING COMPLETION OF ANALYSIS]

---

7. Stakeholder Analysis

7.1 C-Suite Perspectives

CEO / Managing Director:

- Concerns: [List specific concerns for your CEO]
- Success criteria: [What would make this a success in CEO's eyes?]

- Decision factors: [What will drive CEO decision?]

CFO / Finance Director:

- Budget available: [AMOUNT]
- ROI expectations: [CRITERIA]
- Approval threshold: [REQUIREMENTS]

CTO / Technology Director:

- Technical feasibility: [Assessment]
- Engineering capacity: [Available resources]
- Architecture alignment: [Compatibility with current stack]

CISO / Risk Director:

- Compliance priorities: [List]
- Risk reduction targets: [Metrics]
- Audit requirements: [Needs]

Chief Legal Officer / General Counsel:

- Regulatory concerns: [Specific regulations]
- Liability assessment: [Risk areas]
- Due diligence requirements: [Legal needs]

7.2 Operational Teams

Engineering Teams:

- Concerns about implementation complexity: [LIST]
- Required training: [NEEDS]
- Impact on velocity: [ASSESSMENT]

Product Teams:

- Customer-facing implications: [IMPACTS]
- Market positioning: [OPPORTUNITIES]
- Competitive analysis: [DIFFERENTIATION POTENTIAL]

Compliance/Risk Teams:

- Audit support needs: [REQUIREMENTS]
- Documentation requirements: [NEEDS]
- Ongoing monitoring: [CAPABILITIES REQUIRED]

---

9. Measurement & Success Criteria

9.1 Leading Indicators (Months 1-6)

Operational metrics:

- AI decisions requiring human approval: [TARGET %]
- Average human response time: [TARGET]
- System performance overhead: [TARGET]
- Developer satisfaction: [TARGET SCORE]

Track these to validate framework is operating as expected.

9.2 Lagging Indicators (Months 6-24)

Outcome metrics:

- AI-related incidents: [REDUCTION TARGET %]
- Compliance audit findings: [TARGET NUMBER]
- Project success rate: [TARGET %]
- Cost metrics: [ACTUAL vs. PROJECTED]

Track these to validate business case assumptions.

9.3 Qualitative Success Factors

How will you know this was worthwhile?

- [ ] Increased confidence from board/executives
- [ ] Improved customer trust (measured how: [METHOD])
- [ ] Enhanced employee confidence in AI systems
- [ ] Competitive wins attributed to governance
- [ ] Regulatory relationship improvements
- [ ] Industry recognition

---

11. Executive Summary for Decision-Makers

[COMPLETE THIS SECTION LAST, AFTER ALL DATA GATHERED]

The Opportunity

[Describe regulatory/competitive/operational drivers in 2-3 sentences]

Proposed Approach

[Describe Tractatus framework in 2-3 sentences - focus on architectural controls]

Investment Required

- **Total implementation cost:** [AMOUNT]
- **Annual ongoing cost:** [AMOUNT]
- **Timeline:** [DURATION]

Expected Benefits

[List 3-5 primary benefits with evidence/estimates]

Key Risks

[List 3-5 primary risks and mitigations]

Alternatives Considered

[List alternatives and why Tractatus preferred or not]

Recommendation

[APPROVE / DEFER / REJECT] - [Brief rationale]

Next steps: [List immediate actions required]

---

12. Appendices

A. Data Collection Guide

Before completing this template, gather:

From Legal/Compliance:

- ☐ List of applicable regulations
- ☐ Current compliance audit findings

- [ ] Known regulatory risk areas
- [ ] Historical incident reports

From Engineering:

- [ ] Inventory of AI systems in use
- [ ] Technical architecture documentation
- [ ] Integration complexity assessment
- [ ] Engineering capacity availability

From Finance:

- [ ] Budget parameters
- [ ] Cost allocation process
- [ ] ROI calculation methodology
- [ ] Approval thresholds

From Risk Management:

- [ ] Current risk register
- [ ] AI-related incidents/near-misses
- [ ] Risk appetite statement
- [ ] Insurance coverage details

B. Framework Research References

Tractatus Documentation:

- Technical documentation: https://tractatus.sydigital.co.nz/docs.html
- Core concepts: [Link to core concepts doc]
- Implementation guide: [Link to implementer resources]

Framework Status:

- Current status: Research/development framework
- Production deployments: Limited (research implementations)
- Vendor support: SyDigital Ltd (contact@sydigital.co.nz)

Academic Foundations:

- Organizational theory: [Citation]
- AI safety research: [Citation]
- Governance frameworks: [Citation]

C. Regulatory Reference

EU AI Act:

- Official text: Regulation (EU) 2024/1689
- High-risk categories: Annex III
- Compliance timeline: [Key dates]
- Resources: [Links to official sources]

Other Regulations:

- [List sector-specific regulations]
- [Include links to official sources]

D. Decision Log

Use this section to track decision process:

Date	Meeting/Discussion	Attendees	Decisions Made	Next Steps
[DATE]	[MEETING]	[ATTENDEES]	[DECISIONS]	[ACTIONS]

Important Disclaimers

About This Template:

This template is provided as a starting point for organizational assessment. It is not:

- A completed business case ready for presentation
- An assurance of specific outcomes or ROI
- Legal or compliance advice
- A substitute for professional risk assessment
- An endorsement or recommendation of any specific approach

About Tractatus Framework:

The Tractatus Framework is a research/development framework for AI governance. Organizations should:

- Conduct independent technical feasibility assessment
- Validate all claims through pilot testing
- Consult legal counsel for compliance matters

- Obtain vendor quotes for accurate costing
- Assess alternatives appropriate to their context

About Statistical Claims:

Any statistics cited in this template reference industry research (not Tractatus-specific performance). Organizations must:

- Validate applicability to their context
- Measure their own baseline metrics
- Set realistic expectations based on their capabilities
- Avoid extrapolating industry averages to specific situations

Contact: For questions about this template or the Tractatus Framework: contact@sydigital.co.nz

---

This is a template document. It must be completed with organization-specific data before use in decision-making processes.

How to Use This Template

1. **Gather your data** before filling in sections (see Data Collection Guide below)
2. **Replace all [PLACEHOLDER] entries** with your organization's actual information
3. **Delete sections** that don't apply to your situation
4. **Add sections** for organization-specific considerations
5. **Validate assumptions** with relevant stakeholders (Legal, Risk, Finance, Engineering)
6. **Seek expert review** before presenting to decision-makers

⚠️ **Critical:** Do not present this template as a completed analysis. It requires substantial customization based on your organization's reality.

---

2. Tractatus Framework Overview

2.1 What Tractatus Provides

The framework consists of six components designed to create decision boundaries for AI systems:

1. InstructionPersistenceClassifier

- Maintains organizational directives across AI sessions
- Designed to reduce instruction drift over time
- Status: Research implementation, requires adaptation

2. CrossReferenceValidator

- Validates AI actions against established policies
- Designed to detect conflicts before execution
- Status: Research implementation, requires adaptation

3. BoundaryEnforcer

- Prevents AI from making values decisions without human approval
- Designed to preserve human agency for critical choices
- Status: Research implementation, requires adaptation

4. ContextPressureMonitor

- Tracks AI session complexity and token usage
- Designed to detect degraded performance conditions
- Status: Research implementation, requires adaptation

5. MetacognitiveVerifier

- Validates reasoning quality for complex operations
- Designed to improve decision coherence
- Status: Research implementation, requires adaptation

6. PluralisticDeliberationOrchestrator

- Facilitates multi-stakeholder deliberation for values conflicts
- Designed to support non-hierarchical decision-making processes
- Status: Research implementation (October 2025), requires adaptation

2.2 What Tractatus Does NOT Provide

Critical limitations to assess:

- ❌ Not a complete compliance solution (requires integration with broader governance)
- ❌ Not plug-and-play (requires engineering effort to adapt)
- ❌ Not vendor-supported enterprise software (research framework)
- ❌ Not proven at scale in production environments
- ❌ Not a substitute for organizational AI governance processes
- ❌ Not compatible with all AI architectures without modification

Question for your team: Given these limitations, does the architectural approach align with your technical capabilities and risk tolerance?

2.3 Philosophical Foundation

Tractatus is based on the premise that certain decisions are inherently human and should be preserved as such through architectural constraints, not just policy or training.

Core principle: "Whereof the AI cannot safely decide, thereof it must request human judgment."

This differs from approaches that:

- Rely on AI training alone (alignment, RLHF, constitutional AI)
- Use monitoring without structural controls
- Depend on policy enforcement without technical constraints

Assess fit: Does this philosophical approach align with your organization's values and risk management philosophy? □ Yes □ No □ Requires discussion

---

3. Risk Assessment Framework

3.1 Identify Your Risk Categories

For each AI system, assess these risk dimensions:

System	Regulatory Risk	Reputational Risk	Operational Risk	Financial Risk	Total Risk Score
[NAME]	[1-5]	[1-5]	[1-5]	[1-5]	[TOTAL/20]

Risk scoring guidance:

- 1 = Minimal risk
- 2 = Low risk (internal-only, non-critical)
- 3 = Moderate risk (customer-facing, non-high-stakes)
- 4 = High risk (impacts people's lives, regulated decisions)
- 5 = Critical risk (safety-critical, high regulatory exposure)

3.2 Estimate Risk Exposure (Optional)

If you have actuarial or risk modeling capabilities:

For each high-risk system, estimate:

- Probability of adverse event per year: [PERCENTAGE]
- Average cost of adverse event: [AMOUNT]

- Expected annual loss: [CALCULATION]

Note: Most organizations lack sufficient data for accurate estimates. Consider qualitative risk assessment if quantitative data unavailable.

3.3 Current Risk Mitigation

What controls do you currently have?

- □ AI usage policies (policy documents)
- □ Training for AI users
- □ Manual review processes
- □ Access controls
- □ Audit logging
- □ Incident response procedures
- □ Technical controls (specify): [DESCRIPTION]

Gap analysis: What risks remain unmitigated with current controls?

---

5. Benefit Assessment Framework

5.1 Potential Risk Reduction

For each identified risk, estimate potential reduction:

Risk Category	Current Annual Exposure	Estimated Reduction with Tractatus	Residual Risk
Regulatory fines	[AMOUNT or "Unknown"]	[PERCENTAGE]	[AMOUNT]
Reputation damage	[AMOUNT or "Unknown"]	[PERCENTAGE]	[AMOUNT]
Project failures	[AMOUNT or "Unknown"]	[PERCENTAGE]	[AMOUNT]
Compliance costs	[AMOUNT or "Unknown"]	[PERCENTAGE]	[AMOUNT]

⚠️ **Warning:** Estimates should be conservative and validated by risk management professionals. Avoid overstating benefits.

5.2 Operational Efficiency Gains

Where might governance improve efficiency?

- Faster compliance audits: [ESTIMATED HOURS SAVED]
- Reduced rework from AI failures: [ESTIMATED COST AVOIDED]
- Improved project success rates: [ESTIMATED IMPROVEMENT]
- Faster incident response: [ESTIMATED TIME REDUCTION]

Note: These are hypothetical gains. Measure baseline metrics before claiming improvements.

5.3 Strategic Value (Qualitative)

Potential strategic benefits (not quantifiable):

- □ Competitive differentiation through responsible AI
- □ Enhanced customer trust
- □ Improved employee confidence in AI systems
- □ Foundation for future AI initiatives
- □ Regulatory relationship building
- □ Thought leadership opportunities

Question: Which of these matter most to your organization's strategy?

---

10. Risk & Contingency Planning

10.1 Implementation Risks

Risk	Probability	Impact	Mitigation Strategy
Technical integration failure	[H/M/L]	[H/M/L]	[MITIGATION]
Cost overruns	[H/M/L]	[H/M/L]	[MITIGATION]
Timeline delays	[H/M/L]	[H/M/L]	[MITIGATION]
Organizational resistance	[H/M/L]	[H/M/L]	[MITIGATION]
Performance degradation	[H/M/L]	[H/M/L]	[MITIGATION]
Vendor/support issues	[H/M/L]	[H/M/L]	[MITIGATION]

10.2 Contingency Plans

If pilot fails:

- [ ] Rollback plan: [DESCRIPTION]
- [ ] Alternative approach: [ALTERNATIVE]
- [ ] Lessons learned process: [PROCESS]

If costs exceed budget:

- [ ] Scope reduction options: [OPTIONS]
- [ ] Additional funding sources: [SOURCES]
- [ ] Pause criteria: [CRITERIA]

If benefits don't materialize:

- [ ] Measurement review: [PROCESS]
- [ ] Assumption validation: [PROCESS]
- [ ] Continue/abandon decision criteria: [CRITERIA]

Tractatus Framework Assessment Guide

Document Purpose: This template helps organizations evaluate AI governance needs and assess whether the Tractatus Framework approach aligns with their strategic requirements. It is designed to be completed with your organization's actual data, not used as-is.

What This Is NOT: This is not a complete business case with projected ROI figures. Organizations must conduct their own analysis based on their specific risk profile, regulatory exposure, and AI deployment plans.

---

1. Organizational Context Assessment

1.1 Current AI Usage Inventory

Complete this section before proceeding:

System/Tool	Department	Use Case	Data Sensitivity	Regulatory Classification
[NAME]	[DEPT]	[PURPOSE]	[High/Medium/Low]	[EU AI Act category if applicable]
[NAME]	[DEPT]	[PURPOSE]	[High/Medium/Low]	[EU AI Act category if applicable]

Assessment Questions:

- Do you know all AI systems currently in use across your organization? □ Yes □ No □ Uncertain
- Have you identified shadow AI usage (personal accounts for work tasks)? □ Yes □ No □ Uncertain
- Do you know which systems involve customer data or high-stakes decisions? □ Yes □ No □ Uncertain

1.2 Regulatory Exposure

EU AI Act (if applicable):

The EU AI Act establishes penalties for non-compliance:

- Prohibited AI practices: Up to €35M or 7% of global annual turnover (whichever is higher)
- High-risk system violations: Up to €15M or 3% of global annual turnover
- Documentation violations: Up to €7.5M or 1.5% of global annual turnover

Your organization's exposure:

- Annual revenue: [AMOUNT] → Maximum theoretical fine: [CALCULATION]
- Systems classified as high-risk under Annex III: [NUMBER]
- Geographic scope: [Countries where AI systems operate]

Other applicable regulations:

- [List sector-specific regulations: financial, healthcare, employment, etc.]
- [Note: Consult legal counsel for authoritative regulatory analysis]

1.3 Known Incidents & Near-Misses

Historical AI issues in your organization:

Date	Incident Type	Impact	Root Cause	Cost (if known)
[DATE]	[TYPE]	[IMPACT]	[CAUSE]	[COST or "Unknown"]

Industry benchmark: Research indicates 42% of enterprises abandoned AI projects in 2024-2025 due to unclear value and governance challenges. How does your success rate compare?

- Your AI project success rate: [PERCENTAGE or "Unknown"]
- Projects abandoned due to governance concerns: [NUMBER or "Unknown"]

---

4. Implementation Considerations

4.1 Technical Feasibility Assessment

Prerequisites for Tractatus adoption:

Engineering capability:

- Do you have engineers capable of adapting research frameworks? □ Yes □ No
- Estimated engineering capacity available: [NUMBER] engineers for [DURATION]
- Experience with LLM integration: □ Extensive □ Moderate □ Limited □ None

Infrastructure:

- Current LLM providers: [List: OpenAI, Anthropic, internal models, etc.]
- Deployment environment: [Cloud/On-premise/Hybrid]
- Integration complexity: [Simple/Moderate/Complex]

Timeline reality check:

- Research framework adaptation: [ESTIMATED MONTHS]
- Testing and validation: [ESTIMATED MONTHS]
- Production deployment: [ESTIMATED MONTHS]
- **Total estimated timeline:** [TOTAL MONTHS]

4.2 Organizational Readiness

Change management assessment:

- Executive sponsorship secured: □ Yes □ No □ In progress
- Budget authority identified: □ Yes □ No
- Cross-functional team available: □ Yes □ No
- Cultural readiness for AI governance: □ High □ Moderate □ Low

Potential resistance points:

- [List departments/roles that may resist governance controls]

- [List concerns about AI productivity impact]
- [List competing priorities]

4.3 Cost Structure Template

Implementation costs (customize based on vendor quotes):

Phase	Activity	Estimated Cost	Confidence Level
Discovery	Requirements analysis, architecture design	[AMOUNT]	[High/Medium/Low]
Development	Framework adaptation, integration	[AMOUNT]	[High/Medium/Low]
Testing	Validation, security review	[AMOUNT]	[High/Medium/Low]
Deployment	Production rollout, training	[AMOUNT]	[High/Medium/Low]
Total Implementation		**[TOTAL]**	

Ongoing costs (annual):

- Maintenance and updates: [AMOUNT]
- Monitoring and support: [AMOUNT]
- Compliance review: [AMOUNT]
- **Total Annual:** [TOTAL]

Note: These are placeholder estimates. Obtain vendor quotes and internal engineering estimates before presenting financial analysis.

---

6. Alternative Approaches

6.1 Build In-House

Pros:

- Fully customized to organizational needs
- Complete control over architecture
- No vendor dependency

Cons:

- High development cost: [ESTIMATED RANGE]
- Long time to value: [ESTIMATED MONTHS]
- Requires specialized AI safety expertise
- Unproven architecture risk

Estimated cost: [AMOUNT] over [TIMEFRAME]

6.2 Commercial Governance Platforms

Examples: Credo AI, Arthur AI, Fiddler AI, etc.

Pros:

- Vendor-supported enterprise software
- Proven in production
- Compliance reporting built-in

Cons:

- Monitoring focus, not architectural controls
- SaaS pricing can be high
- May not address decision boundary concerns

Estimated cost: [AMOUNT] annual subscription

6.3 Consulting-Led Frameworks

Examples: McKinsey, Deloitte, PwC AI governance consulting

Pros:

- Comprehensive governance approach
- Strong compliance coverage
- Executive-level engagement

Cons:

- Policy-based, not technical enforcement
- High consulting fees
- Requires ongoing organizational discipline

Estimated cost: [AMOUNT] for [DELIVERABLES]

6.4 Do Nothing / Maintain Current State

Pros:

- Zero additional investment
- No organizational disruption

Cons:

- Regulatory risk exposure continues
- Competitive disadvantage as others adopt governance
- Potential for costly incidents

Estimated cost: [CURRENT RISK EXPOSURE]

6.5 Tractatus Framework Adaptation

Pros:

- Architectural approach to decision boundaries
- Research framework with documented approach
- Open for organizational adaptation

Cons:

- Research-stage, not mature commercial product
- Requires engineering investment to adapt
- Limited vendor support
- Unproven at enterprise scale

Estimated cost: [AMOUNT for implementation + adaptation]

Decision criteria: Which approach best balances your technical capability, risk tolerance, and budget constraints?

---

8. Decision Framework

8.1 Go/No-Go Criteria

Must-Have Requirements:

- □ Executive sponsorship secured
- □ Budget approved: [AMOUNT]
- □ Engineering capacity allocated

- □ Regulatory driver confirmed
- □ Technical feasibility validated

Should-Have Requirements:

- □ Cross-functional team committed
- □ Pilot use case identified
- □ Success metrics defined
- □ Change management plan developed

Nice-to-Have:

- □ Industry peer validation
- □ Customer interest confirmed
- □ Competitive intelligence supports decision

Decision: Proceed if [NUMBER] of Must-Have + [NUMBER] of Should-Have criteria met.

8.2 Recommended Next Steps

If proceeding:

1. **Month 1:**
    - ☐ Assign executive sponsor
    - ☐ Form cross-functional team
    - ☐ Engage vendor for detailed scoping
    - ☐ Identify pilot system(s)
2. **Month 2-3:**
    - ☐ Complete technical feasibility study
    - ☐ Develop detailed implementation plan
    - ☐ Secure final budget approval
    - ☐ Initiate procurement process
3. **Month 4+:**
    - ☐ Begin framework adaptation
    - ☐ Pilot deployment
    - ☐ Measure and validate

If not proceeding:

- ☐ Document decision rationale

- [ ] Revisit in [TIMEFRAME]
- [ ] Pursue alternative: [SELECTED ALTERNATIVE]

Document Control

Version: 2.0 (Template version) **Last Updated:** 2025-10-09 **Document Type:** Template - Requires Completion **Classification:** Internal Use - Customize Before External Distribution **Owner:** [ASSIGN DOCUMENT OWNER]

Completion Status:

- [ ] Data collection complete
- [ ] All placeholders replaced
- [ ] Financial analysis validated
- [ ] Risk assessment completed
- [ ] Stakeholder input gathered
- [ ] Legal review completed
- [ ] Executive summary drafted
- [ ] Ready for decision-maker presentation

Next Review: [DATE]

Document Metadata

- **Version:** 2.0
- **Created:** 2025-10-09
- **Last Modified:** 2025-10-13
- **Author:** SyDigital Ltd
- **Word Count:** 3,250 words
- **Reading Time:** ~16 minutes
- **Document ID:** business-case-tractatus-framework
- **Status:** Active (Template)

License

Additional Terms:

1. **Attribution Requirement**: Any use, modification, or distribution of this work must include clear attribution to the original author and the Tractatus Framework project.
2. **Moral Rights**: The author retains moral rights to the work, including the right to be identified as the author and to object to derogatory treatment of the work.
3. **Research and Educational Use**: This work is intended for research, educational, and practical implementation purposes. Commercial use is permitted under the terms of the Apache 2.0 license.
4. **No Warranty**: This work is provided "as is" without warranty of any kind, express or implied. The author assumes no liability for any damages arising from its use.
5. **Community Contributions**: Contributions to this work are welcome and should be submitted under the same Apache 2.0 license terms.

For questions about licensing, please contact the author through the project repository.